

Accelerating Science Through Applied Research Computing

**Community for Data Integration
10 September 2014**

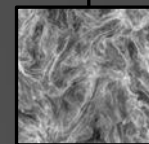
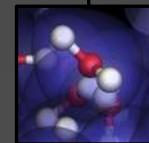
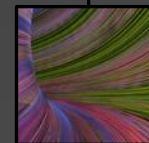
Jeff Falgout

jfalgout@usgs.gov

US Geological Survey

Core Science Analytics, Synthesis & Libraries

Applied Research Computing



Agenda

- ✧ What is Applied Research Computing
- ✧ Overview of some research computing use cases
- ✧ FY15 and beyond roadmap for Applied Research Computing

Applied Research

“Applied Research is also original investigation undertaken in order to acquire new knowledge. It is, however, directed primarily towards a specific practical aim or objective.”

OECD (2002), Frascati Manual 2002: Proposed Standard Practice for Surveys on Research and Experimental Development, The Measurement of Scientific and Technological Activities, OECD Publishing.
DOI: 10.1787/9789264199040-en

Research Computing

“Research Computing refers to the advanced computing resources, software, hardware, and personnel that are required by researchers from any discipline.”

<http://www.usf.edu/it/research-computing/what-is-rc/index.aspx>

Applied Research Computing

Applied Research Computing is investigation undertaken in order to acquire new knowledge directed primarily towards advanced computing resources, software, hardware, and personnel, that are required by researchers from any discipline.

Applied Research Computing is **NOT**

- ✧ Enterprise IT

 - ✧ Email

 - ✧ Helpdesk

 - ✧ Desktop Support

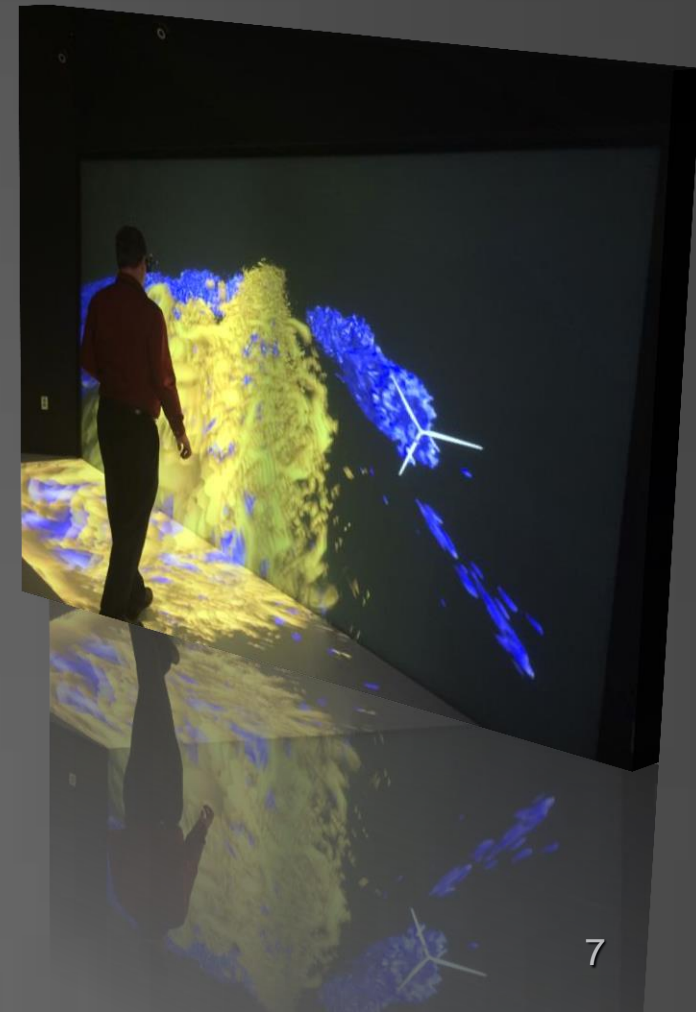
- ✧ Standard Development Methods

 - ✧ Works towards a defined goal or product

 - ✧ Predefined requirements that development team must meet

Applied Research Computing IS

- ✧ A wandering path to discovery
- ✧ Final outcomes unknown
- ✧ Can be messy and confusing
- ✧ Hard to predict timelines
- ✧ Iterative development – let the data be your guide
- ✧ Enormous, cumbersome data volumes
 - ✧ Big Data??



BIG DATA!!!

“Big data is like teenage sex: everyone talks about it, nobody really knows how to do it, everyone thinks everyone else is doing it, so everyone claims they are doing it ...”

(Dan Ariely)



A Tipping Point ...

- ✧ At some point in research there is a need to:
 - ✧ Expand beyond our current study area
 - ✧ Regional, National, Global
 - ✧ Integrate a new data layer
 - ✧ Increase the model resolution
 - ✧ Decrease processing time
- ✧ These changes result in rendering current methods less than useful
 - ✧ Processing on my desktop or single server no longer works well ...

Beyond the Desktop

- ✧ What to do when the desktop machine isn't enough? Upgrade to a faster processor!
 - ✧ Not any more ... CPU speeds have leveled off
 - ✧ Instead of faster CPUs, machines are increasing the NUMBER of CPUs

Enter ...

High Performance Research Computing



- ✧ Reduce the time to solution
- ✧ Increase the scope of the question

CSAS's Applied Research Computing

- ✧ In FY2013, Core Science Analytics and Synthesis, under the CSS mission area, formed a pilot program to provide a resource for research scientists needing assistance with research/scientific computing
- ✧ As of now (September 2014), we have a budget, hardware, and five full time staff devoted to the program. We have moved beyond the pilot stage.

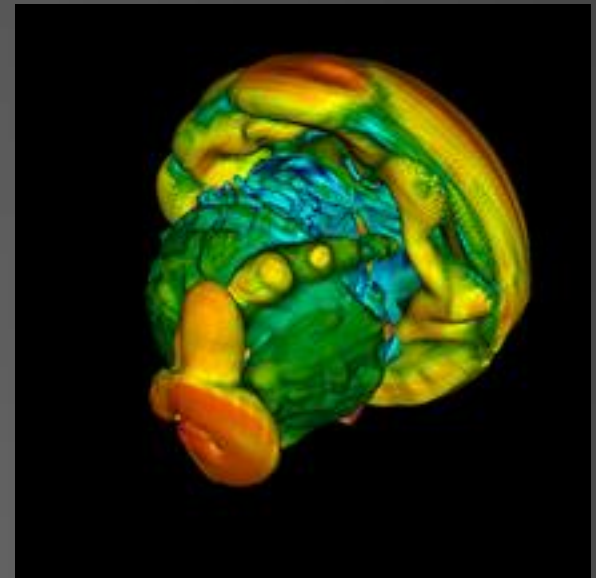
Applied Research Computing

- ✧ ... is a program (little p) devoted to assisting USGS research staff in the area of research computing
 - ✧ No cost to the project or PI
 - ✧ Except for cases of extensive development work (hasn't happened yet)
 - ✧ All that is required is an Active Directory account for access to ARC's hardware resources

Applied Research Computing's Goals

Provide scientific high performance computing (HPC), high performance storage (HPS), high capacity storage (HCS) expertise, education, and resources to scientists, researchers and collaborators.

- ✧ Decrease “time to solution”
 - ✧ Faster results
- ✧ Increase “scope of question”
 - ✧ Bigger question
 - ✧ Higher accuracy



Volume rendering of a supernova
generated by Jaguar

Applied Research Computing – Hardware/Software

- ✧ ~700 Core distributed memory cluster
- ✧ Several Compilers
- ✧ Several MPI libraries
- ✧ Aim to track what most HPC centers provide
 - ✧ Minimize headaches when moving to large resources



A Sampler of Current Projects in the Work Queue of Applied Research Computing

Modeling of Watershed Systems

Lauren Hay, Steve Markstrom, Steve Reagan, Roland Viger, et al. NRP

- ✧ Access to computing cluster
 - ✧ “Changed the way we do business”
- ✧ Introduced parallel processing techniques
 - ✧ ARC Tutorial and support
 - ✧ CU workshop
- ✧ Currently scoping a code optimization service to increase performance and introduce parallelism in models

Modeling species response to environmental change: development of integrated, scalable Bayesian models of population persistence - Ben Letcher, Leetown SC

- ✧ R, Bugs, and Jags
- ✧ Running MCMC chains in parallel
- ✧ Large memory footprint
 - ✧ Up to 12G per process
- ✧ Extremely long run times on some rivers
 - ✧ Mississippi River models are at 159 days and counting
- ✧ Lesson in implementing checkpoint/restart
 - ✧ Failures in hardware have interrupted jobs – months of processing time lost

Burned Area Mapping Algorithm

Essential Climate Variables – Todd Hawbaker, GECSC

- ✧ First prototype in R
 - ✧ Scaling issues with processing 14TB Landsat Imagery
 - ✧ Single scene runtime ~14 days
- ✧ Converted to Python, Implemented MPI and parallel processing strategies
 - Single scene runtime ~80 minutes and shrinking
 - ~26 days to run all 459 scenes



Caribou Gene Flow

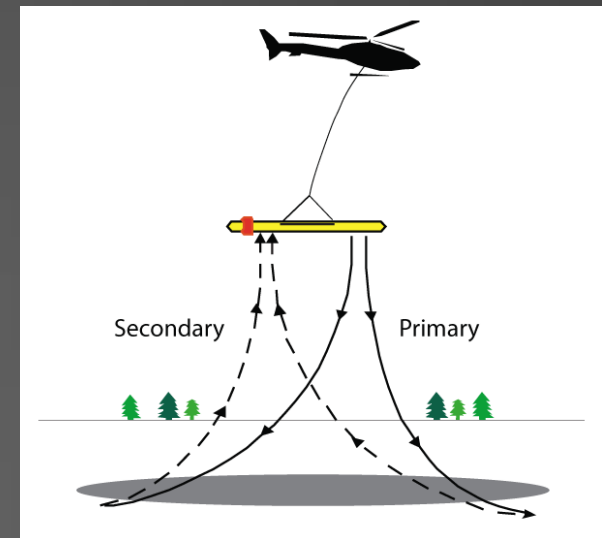
Laura Thompson - NCCWSC

- ✧ Circuitscape used to model population gene flow (Python based)
- ✧ PI provides input files and variables
- ✧ We launch jobs and return output files
 - ✧ Infrequent jobs, jointly determined that this was best path

Airborne Electro-Magnetic Surveying

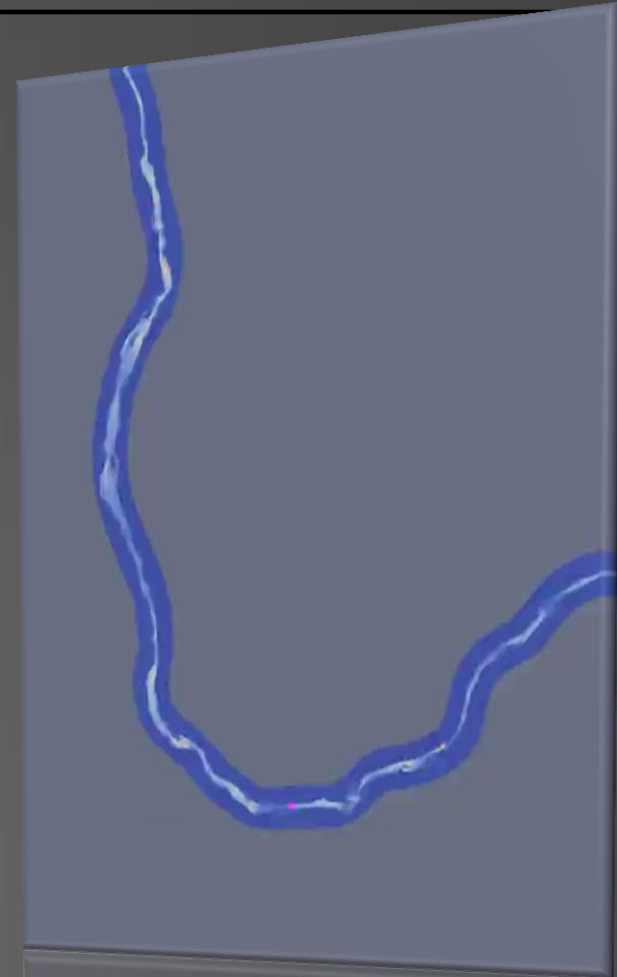
Burke Minsley, Andy Kass, Paul Bedrosian - CGGSC

- ✧ Debug ModularEM MPI issues, provide a fix to upstream maintainers (Fortran code)
- ✧ Additional Computing Platform
 - ✧ Saving money by not replacing their current cluster
 - ✧ Migrate most of computing work to CSAS cluster
- ✧ MatLab to Python translation
 - ✧ 1000's of MC simulations



Sturgeon Larvae Drift in the Snake River of Idaho – Rich McDonald, NRP

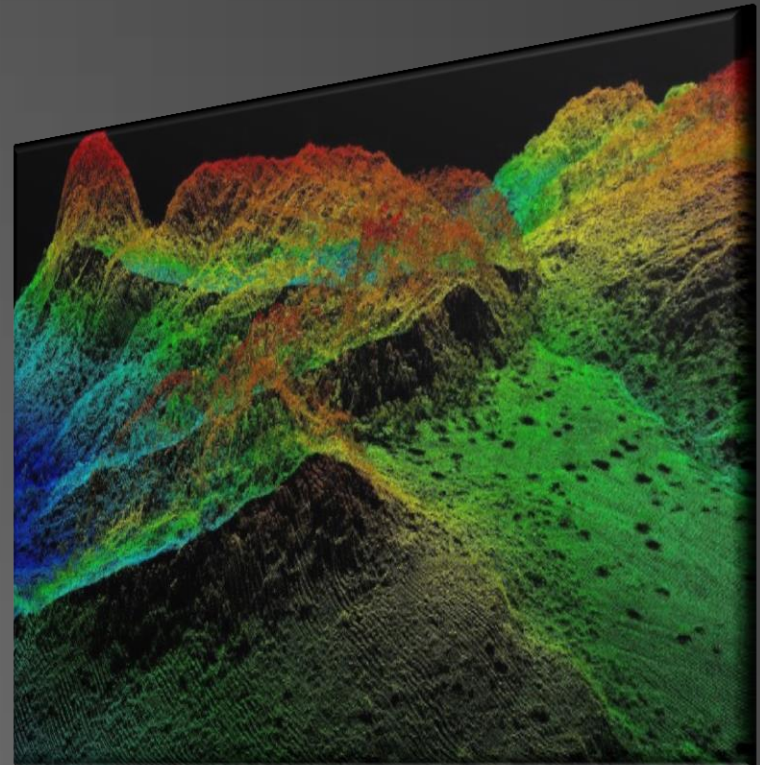
- ✧ Particle tracking simulation
- ✧ Issues with scaling number of particles per simulation
 - ✧ 5,000 particles over 24 hours took 36 hours to compute
 - ✧ Slower than real time!!!
- ✧ Actively working on performance improvements
 - ✧ 10,000 particles over 24 hours now takes 3-4 hours to compute
 - ✧ Still working on performance
- ✧ Introducing code based checkpoints and restarting



LiDAR in the Great Smoky Mountains NP

John Kosovich, Jeff Wendel CSAS

- ✧ Currently using Global Mapper, ArcGIS to process point clouds
 - ✧ Long processing times
 - ✧ Licensed software
 - ✧ Not scalable
- ✧ Developing new methods for add-on Products to LiDAR data
 - ✧ Software for HPC environment to process data at larger scale



Numbers to Date

- ✧ Currently working with ~20 science projects
 - ✧ Several have completed work with ARC computing resources and are in “manuscript mode”
- ✧ Launched over 24,150 compute jobs
- ✧ Served ~1,000,000 CPU hours and counting

ARC/CSAS Funded Science Center Projects

- ✧ Essential Climate Variables Burned Area Mapping Algorithm - GECSC
 - ✧ Landsat imagery analysis
 - ✧ HPC storage upgrade
 - ✧ High I/O
- ✧ Cascades/Alaska Volcano Observatories
 - ✧ Matlab license
 - ✧ GPGPU upgrade



- ✧ Agreement with Research Computing at CU
 - ✧ <http://www.rc.colorado.edu>
- ✧ Five Year Agreement, NTE \$100,000
- ✧ Focus on research/scientific computing education/training for USGS research staff
- ✧ July 2014 initial workshop on DFC
 - ✧ Topics included intro to Linux, HPC101, Python, OpenMP, MPI

Oak Ridge National Lab (ORNL)

Interagency Agreement

- ✧ We have on-demand access to the Environmental Sciences Division workgroup clusters
 - ✧ 832 core and 256 core clusters
- ✧ Commitment from the ORNL Leadership Computing Facility director to entertain director's discretion allocations on the supercomputer Titan
- ✧ Several staff with HPC expertise have agreed to assist USGS when requested

Rocky Mountain Advanced Computing Consortium (RMAcc)



- ✧ Regional group with members from universities and government
 - ✧ NOAA, NREL, INL, NCAR, USGS
 - ✧ U of CO, CO State, CO School of Mines, U of WY, U of ID, ID State, Boise State, U of Utah, Montana State, U of NM
- ✧ www.rmacc.org
- ✧ HPC Symposium Aug 12-13th, 2014
 - ✧ 2015 Symposium Date/Location TBD

XSEDE Champions



Extreme Science and Engineering
Discovery Environment

- ✧ Local representative and liaison between XSEDE and campus (or organization)
- ✧ Local source of knowledge concerning HPC and other digital resources available through XSEDE
- ✧ MOU signed by Kevin Gallagher May 2014 establishing two USGS Champions
 - ✧ Jeff Falgout – jfalgout@usgs.gov
 - ✧ Janice Gordon – janicegordon@usgs.gov

What is XSEDE?

- ✧ The eXtreme Science and Engineering Discovery Environment <http://www.xsede.org/>
- ✧ Five year, \$121 Million project supported by the National Science Foundation
- ✧ National Cyberinfrastructure for U.S. based research institutions
- ✧ Replaces and expands TeraGrid
- ✧ Single virtual system for interactive sharing of computing resources, data, and expertise
 - ✧ Initially 16 supercomputers and high-end visualization systems and now includes other specialized digital resources.
 - ✧ No cost to scientists

XSEDE's Goal



Extreme Science and Engineering
Discovery Environment

- ✧ Lower the technical barriers to the access and use of computing resources
- ✧ Establishes private, secure environments that have all the resources, services, and collaboration support for researchers to be productive

- ✧ SGI UV 2000 Shared Memory, Single System Image Machine
 - ✧ 256 Core (32 CPU) Intel Ivy Bridge E5-4627v2 3.3 GHz
 - ✧ 4 TB DDR3 Main Memory
 - ✧ Globally Addressable, Cache Coherent
 - ✧ NUMalink6 node interconnect
 - ✧ 6GB/s ultra low latency
 - ✧ 54 TB Global Scratch
 - ✧ 16 Gb/s Fiber Channel
 - ✧ CXFS file system
 - ✧ 80 TB project storage



- ✧ FY15 Planned Upgrades
 - ✧ At least an additional 256 cores
 - ✧ Additional 32 CPUs
 - ✧ Intel Xeon Phi
 - ✧ Nvidia Tesla K20 GPUs
 - ✧ Additional 4 TB Main Memory
 - ✧ 240-480 TB additional storage



Lustre Research

- ✧ Distributed parallel file system
- ✧ Designed for high IO bandwidth applications
 - ✧ File system of choice for many supercomputers
- ✧ Planned FY15 prototype/testing of lustre for distributed memory nodes
 - ✧ Could lead to FY16 hardware upgrade

Software Upgrades

✧ Matlab

✧ Distributed Computing Server

✧ 192 Workers

✧ Parallel Computing Toolbox

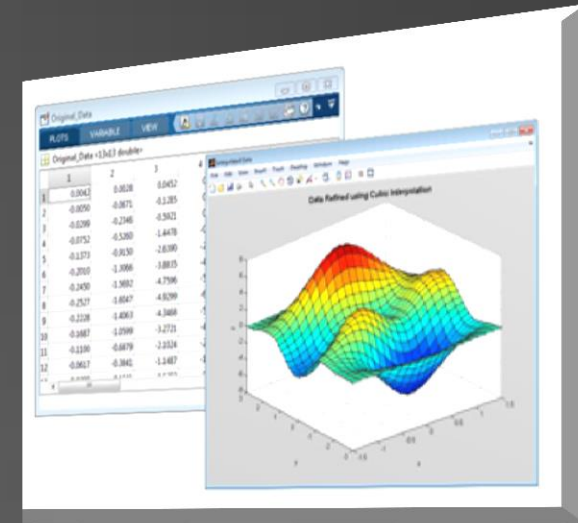
✧ Two floating licenses

✧ RFQ Closed 9/8/2014 – award this week

✧ Combined with Yeti, possibility to run 448 workers at once (needs confirmation).

✧ TotalView Parallel Debugger

✧ FY15 Budgeted



Training Workshops for FY15



- ✧ Two workshops through the CU CESU
 - ✧ Jan/Feb 2015
 - ✧ Data Analysis, Visualization, Movement, Management
 - ✧ Summer 2015
 - ✧ Advanced HPC topics including
- ✧ SGI UV Training
 - ✧ New processing paradigm with 4TB RAM

Student Development program

- ✧ FY15 partnership with AEM/CGGSC
- ✧ ARC program will hire, mentor, and provide office space for student with hpc focus and, hopefully, geophysics background
- ✧ AEM/CGCSC will fund student and provide science problems

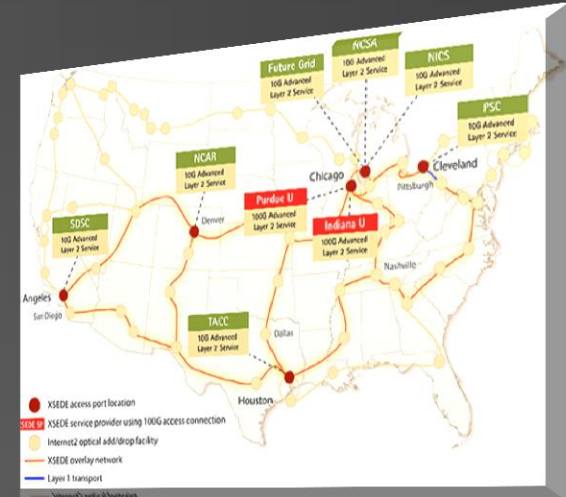
HPC Owners CO-OP

Planned (delayed from FY13) endeavor to establish a coordinating cooperative among the owners of HPC resources within USGS

- ✧ Establish methods for USGS researchers to access resources within USGS
- ✧ Plan upgrades
- ✧ Share expertise

Science DMZ and Data Transfers

- ✧ <http://fasterdata.es.net/science-dmz/>
- ✧ Research grade networking
 - ✧ 10-100 GB/s WAN
 - ✧ Internet2 Consortium
 - ✧ ES.net
 - ✧ BiSON (Bi-State Optical Network) (CO, WY)
- ✧ Data transfer methods
 - ✧ GridFTP, Globus
 - ✧ Establishment of Data Transfer Nodes
 - ✧ Huge volumes of data present new problems ...



Questions? Comments?

- ✧ ARC Help – gs-css_csas_hpc_help@usgs.gov
- ✧ Jeff Falgout – jfalgout@usgs.gov
- ✧ Janice Gordon – janicegordon@usgs.gov
 - ✧ Note there are two Janice Gordons in the USGS ...

